Séminaire BOURBAKI Mars 2021
72e année, 2019–2021, no 1174, p. 237 à 273
doi : 10.24033/ast.1163

QUANTITATIVE INVERSE THEORY OF GOWERS UNIFORMITY NORMS [after F. Manners]

by **Thomas F. Bloom**

INTRODUCTION

Let $s \geq 1$ be a fixed integer (in particular, all implicit constants may depend on s). Fix some large prime N and use $\mathbb{Z}_N$ to denote the cyclic group $\mathbb{Z}/N\mathbb{Z}$. If $f : \mathbb{Z}_N \to \mathbb{C}$ and $h \in \mathbb{Z}_N$ we define the multiplicative derivative

$$\Delta_h f : \mathbb{Z}_N \longrightarrow \mathbb{C}, \quad \Delta_h f(x) = f(x)\overline{f(x+h)}.$$

This definition has a natural extension to vectors $\boldsymbol{h} \in \mathbb{Z}_N^k$, so that

$$\Delta_{\boldsymbol{h}} f(x) = \Delta_{h_1} \cdots \Delta_{h_k} f.$$

We can now define our central object of study: the Gowers uniformity norm of degree s. This is defined for $f : \mathbb{Z}_N \to \mathbb{C}$ by

$$\|f\|_{U^{s+1}} = \Big(\mathop{\mathbb{E}}_{\boldsymbol{h} \in \mathbb{Z}_N^{s+1}} \mathop{\mathbb{E}}_{x \in \mathbb{Z}_N} \Delta_{\boldsymbol{h}} f(x) \Big)^{1/2^{s+1}}.$$

Here we have used the expectation notation, now standard within additive combinatorics, to denote the normalised sum (so that, for example, $\mathbb{E}_{x\in H}$ means $\frac{1}{|H|}\sum_{x \in H}$). Alternatively, one could use the inductive definition

$$\|f\|_{U^1} = \Big| \mathop{\mathbb{E}}_{x} f(x) \Big| \quad \text{and} \quad \|f\|_{U^{s+1}} = \Big(\mathop{\mathbb{E}}_{h \in \mathbb{Z}_N} \|\Delta_h f\|_{U^s}^{2^s} \Big)^{1/2^{s+1}}.$$

These norms [1] were introduced into additive combinatorics by Gowers (2001) in his analytic proof of Szemerédi's theorem. They provide a quantitative measure of the additive structure of f. For example, the U^2 norm

$$\|f\|_{U^2} = \Big(\mathop{\mathbb{E}}_{x, h_1, h_2 \in \mathbb{Z}_N} f(x)\overline{f(x+h_1)f(x+h_2)}f(x+h_1+h_2) \Big)^{1/4}$$

1. It is not obvious from the definition that they are indeed norms, but this is true for $s \geq 1$. A proof can be found in the original paper of Gowers (2001).

measures f along 2-dimensional additive quadruples of the shape

$$(x, x + h_1, x + h_2, x + h_1 + h_2).$$

In general, the U^{s+1} norm measures f along 2^{s+1}-tuples which are (the projections of) $(s + 1)$-dimensional cubes.

These configurations are generic enough that if we can understand how f behaves on them then we can understand f on many other kinds of linear system. For example, the count of f along arithmetic progressions of length k is strongly related to $\|f\|_{U^{k-1}}$. As a result the uniformity norms play a central role in modern additive combinatorics. They often appear in proofs which use a 'structure vs. randomness' philosophy. Such proofs go along the following lines:

1) Find some function f such that the behavior of the objects one wishes to understand is governed by some $\|f\|_{U^{s+1}}$ (for example, when counting k-term arithmetic progressions inside some $A \subset \mathbb{Z}_N$ one would take $f = 1_A - |A|/N$ and $s = k - 2$).

2) Show that the theorem in question follows if $\|f\|_{U^{s+1}}$ is *small* – this is the random aspect of the 'structure vs. randomness' dichotomy, and usually follows from elementary counting methods.

3) Show that if $\|f\|_{U^{s+1}}$ is *large* then f has some algebraic structure.

4) Finally conclude the proof by showing that if f has this algebraic structure then the theorem also follows.

Landmark applications of this method are the analytic proof of Szemerédi's theorem by Gowers (2001), the first to deliver reasonable quantitative bounds, and the proof by Green and Tao (2008b) that the primes contain arbitrarily long arithmetic progressions.

The inverse theory of uniformity norms, which is the focus of this article, addresses the third point, and seeks to answer the following question.

Question (The Inverse Question). *If a* 1*-bounded* [2] *function* $f : \mathbb{Z}_N \to \mathbb{C}$ *has large* U^{s+1} *norm then what can we deduce about* f*?*

We will first explore this question in the simplest case $s = 1$. A brief experimentation shows that the function [3] $x \mapsto e(\alpha x)$, where $\alpha \in \frac{1}{N}\mathbb{Z}$ (so that this does indeed give a well-defined function on $\mathbb{Z}_N$), satisfies $\|f\|_{U^2} = 1$, the maximum possible. Therefore (linear) characters provide examples of bounded functions with large U^2 norm. The inverse theorem for the U^2 norm says that the converse is also true, at least in an approximate sense. More precisely, if $\|f\|_{U^2} \geq \delta$ then f must correlate with a linear character, in that there exists some $\alpha \in \frac{1}{N}\mathbb{Z}$ such that

$$\left|\mathop{\mathbb{E}}_{x} f(x)\overline{e(\alpha x)}\right| > c(\delta) \tag{1}$$

2. A function f is 1-bounded if $|f(x)| \leq 1$ for all x.
3. We use $e(x)$ to denote $\mathrm{e}^{2\pi i x}$.

for some constant $c(\delta) > 0$ depending only on δ. The proof of this inverse result is very short. We note first that the U^2 system of quadruples of the shape

$$(x, x + h_1, x + h_2, x + h_1 + h_2)$$

are exactly those quadruples (x_1, x_2, x_3, x_4) such that $x_1 - x_2 = x_3 - x_4$. We may use linear characters and orthogonality to detect this linear condition so that, defining the Fourier transform by

$$\widehat{f}(a) = \mathop{\mathbb{E}}_{x} f(x)e(-ax/N)$$

for $a \in \mathbb{Z}_N$, we have

$$\begin{aligned}\|f\|_{U^2}^4 &= \mathop{\mathbb{E}}_{x_1 - x_2 = x_3 - x_4} f(x_1)\overline{f(x_2)f(x_3)}f(x_4)\\ &= \sum_{a \in \mathbb{Z}_N} \left|\widehat{f}(a)\right|^4 \le \sup_a \left|\widehat{f}(a)\right|^2 \sum_{a \in \mathbb{Z}_N} \left|\widehat{f}(a)\right|^2\\ &= \sup_a \left|\widehat{f}(a)\right|^2 \mathop{\mathbb{E}}_{x \in \mathbb{Z}_N} \left|f(x)\right|^2 \le \sup_a \left|\widehat{f}(a)\right|^2,\end{aligned}$$

using Parseval's identity and the assumption that f is 1-bounded. It follows that if $\|f\|_{U^2} \ge \delta$ then (1) holds with the explicit lower bound $\ge \delta^2$.

For higher uniformity norms with $s \ge 2$ this simple argument fails, and it is less clear what functions have large U^{s+1} norm. We note that the uniformity norms are nested, in the sense that

$$\|f\|_{U^2} \le \|f\|_{U^3} \le \|f\|_{U^4} \le \cdots,$$

and so in particular this inverse question becomes more difficult as s increases, since any example of a function with large U^{s+1} norm certainly also has large U^{s+2} norm, but the converse may not hold.

Considering what happens when $s = 1$, we observe that the reason exponentials with linear phases have large U^2 norm is because the second derivative of a linear function always vanishes (hence $\Delta_{h_1,h_2} f$ is an exponential with phase 0, and so is identically 1). Generalising this observation shows that $f(x) = e(P(x))$ will have $\|f\|_{U^{s+1}} = 1$ whenever P is a polynomial of degree $\le s$ (with coefficients in $\frac{1}{N}\mathbb{Z}$ so that it is well-defined on $\mathbb{Z}_N$).

At this point one may guess that, just as when $s = 1$, the approximate converse also holds, and conjecture something like: if $\|f\|_{U^{s+1}} \ge \delta$ then

$$\left|\mathop{\mathbb{E}}_{x} f(x)\overline{e(P(x))}\right| > c(\delta) \tag{2}$$

for some polynomial $P(x) \in \frac{1}{N}\mathbb{Z}[x]$ of degree $\le s$ and some constant $c(\delta) > 0$ depending only on δ. It turns out this is not quite enough, and the exponentials with polynomial phases do not represent a full set of obstructions. [4] To see why, observe

4. The term obstructions here comes from thinking of functions g such that "$|\mathbb{E}_x f(x)\overline{g(x)}| \gg 1$ implies $\|f\|_{U^{s+1}} \gg 1$" as 'obstructions' to having small uniformity norm.

that we don't require our 'phase functions' $P(x)$ to actually vanish after taking $s+1$ derivatives. For the U^{s+1} norm of $e(P(x))$ to be large, it suffices for the derivative of $P(x)$ to be biased towards 0 (so that the multiplicative derivative of $e(P(x))$ is biased towards 1). Consider, for example, the function on $\mathbb{Z}_N$ defined by [5]

$$x \longmapsto \{\alpha x\}\{\beta x\}$$

for $\alpha, \beta \in \frac{1}{N}\mathbb{Z}$. This is not a true quadratic, since the third derivative is not identically zero. This failure arises because $x \mapsto \{\alpha x\}$ is not truly linear. We note, however, that it is linear a large proportion of the time, since we can write

$$\{\alpha(x+y)\} = \{\alpha x\} + \{\alpha y\} - \rho_{x,y},$$

where $\rho_{x,y}$ is a 'carry bit' function that is 1 if $\{\alpha x\} + \{\alpha y\} \geq 1$ and 0 otherwise. In particular, $\rho_{x,y} = 0$ with probability $\frac{1}{2}$ and hence $\{\alpha x\}$ behaves like a linear function with probability $\frac{1}{2}$. By a similar calculation, the third derivative of a quadratic function like $P(x) = \{\alpha x\}\{\beta x\}$ vanishes a positive proportion of the time, and thus $e(P(x))$ has large U^3 norm. Thus we need to also include these generalised bracket polynomials (generated by repeated composition of, not just addition and multiplication, but also the fractional part operator $\{\bullet\}$) as possible obstructions.

The inverse theorem for the Gowers uniformity norms states that this expanded set of obstructions captures all the reasons that a function might have large uniformity norm. That is, if $\|f\|_{U^{s+1}} \geq \delta$ then there exists some bracket polynomial P of degree $\leq s$ such that (2) holds. Such a statement was conjectured by Green and Tao (2010b) in their work on linear equations in primes and first proved in a qualitative sense (that is, with no bounds on the function $c(\delta)$) by Green, Tao, and Ziegler (2012).

The focus of this article is a recent new proof of the inverse theorem by Manners (2018) which gives, for the first time, a quantitative version of this statement. Our aims are threefold:

1) to state precisely the inverse theorem as proved in Manners (2018), defining all the concepts required;
2) to sketch some of the ideas used in the proof, giving a flavour of the kind of arguments used, and
3) to state precisely some of the new definitions and concepts introduced by Manners.

The paper of Manners (2018) itself is over 100 pages long, so we will not be able to even approach a proper proof of the inverse theorem. As a result of this intimidating length, however, some of the beautiful new ideas introduced by Manners may go otherwise unnoticed by those without the time to plumb the technical depths. We hope that this article helps popularise them, so that they can find many other applications.

5. As usual in number theory, we use $\{\bullet\} : \mathbb{R} \to [0,1)$ to denote the fractional part operator $x \mapsto x - \lfloor x \rfloor$, where $\lfloor x \rfloor$ is the largest integer n such that $n \leq x$.

Acknowledgements

I would like to thank Freddie Manners for many helpful comments and guidance on understanding his work, and Aled Walker for his notes and insights on Manners' proof, which have helped me immensely. I would also like to thank Nicolas Bourbaki, Gabriel Conant, Ben Green, and Sarah Peluse for their comments on an earlier draft.

1. THE INVERSE THEOREM

The following quantitative inverse theorem for the Gowers uniformity norms is the main result of Manners (2018), and the focus of this article.

Theorem 1.1 (Manners). — *Let $\delta > 0$ and $N \geq 2$ be prime. If $f : \mathbb{Z}_N \to \mathbb{C}$ is a 1-bounded function such that $\|f\|_{U^{s+1}} \geq \delta$ then there exists $\epsilon > 0$ and a 1-bounded, N-periodic, nilsequence $\psi : \mathbb{Z} \to \mathbb{C}$ with degree s, dimension D, parameter K, and complexity M, such that*

$$\Big| \mathop{\mathbb{E}}_{x} f(x)\, \overline{\psi(x)} \Big| \geq \epsilon,$$

where the parameters are bounded in terms of δ by

$$D = O(\delta^{-O(1)}) \quad \textit{and} \quad \epsilon^{-1}, K, M \leq \begin{cases} \exp(O(\delta^{-O(1)})) & \textit{if } s \leq 3 \textit{ and} \\ \exp(\exp(O(\delta^{-O(1)}))) & \textit{if } s \geq 4. \end{cases}$$

We have followed Manners in stating the inverse result using nilsequences rather than bracket polynomials. These are qualitatively equivalent, as shown by Bergelson and Leibman (2007). [6] Nilsequences are usually, however, easier to work with. We defer an explicit definition of nilsequences, along with the meaning of degree, dimension, parameter, and complexity, till Section 3. For now the reader should think of a nilsequence of degree s as a function of the shape $n \mapsto e(P(n))$, where $P(n)$ is a (bracket) polynomial of degree at most s.

Since it can also be shown that if f correlates with a nilsequence of degree s then it has large U^{s+1} norm (we sketch a proof of this in Section 3.4), this gives a complete characterisation of functions with large uniformity norm. A characterisation of this type was first conjectured by Green and Tao (2010b, Conjecture 8.3) in their work on linear equations in primes.

Before we explore Theorem 1.1 and its proof we briefly summarise other, related, results.

6. Bergelson and Leibman (2007) do not give any quantitative form of this equivalence, which would be needed to make a version of Theorem 1.1 precise for bracket polynomials. Since nilsequences are more convenient to work with anyway we will not address this, and use bracket polynomials only as motivational examples.